Online Budgeted Truthful Matching

Rahul Vaze
School of Technology and Computer Science,
Tata Institute of Fundamental Research,
Mumbai, India

Marceau Coupechoux
LTCI, CNRS, Telecom ParisTech, University
Paris-Saclay, France

arXiv:1606.00581v1 [cs.DS] 2 Jun 2016

ABSTRACT

An online truthful budgeted matching problem is considered for a bipartite graph, where the right vertices are available ahead of time, and individual left vertices arrive sequentially. On arrival of a left vertex, its edge utilities (or weights) to all the right vertices and a corresponding cost (or bid) are revealed. If a left vertex is matched to any of the right vertices, then it has to be paid at least as much as its cost. The problem is to match each left vertex instantaneously and irrevocably to any one of the right vertices, if at all, to find the maximum weight matching that is truthful, under a payment budget constraint. Truthfulness condition requires that no left vertex has any incentive of misreporting its cost. Assuming that the vertices arrive in an uniformly random order (secretary model) with arbitrary utilities, a truthful algorithm is proposed that is 24β-competitive (where β is the ratio of the maximum and the minimum utility) and satisfies the payment budget constraint. Direct applications of this problem include crowdsourcing auctions, and matching wireless users to cooperative relays in device-to-device enabled cellular network.

1. INTRODUCTION

Motivated by applications in crowdsourcing and device-to-device (D2D) communication, we consider an online bipartite matching problem over a bipartite graph $G(L \cup R, E)$, where the right vertex set R is known ahead of time, while left vertices of L arrive sequentially in a random order. The incident edge utilities from a vertex $\ell \in L$ to set R are revealed only upon its arrival, as well as its cost c_ℓ, and the problem is to decide which vertex of R to match with ℓ, if at all, immediately and irrevocably. If vertex ℓ is matched, a payment p_ℓ is made to vertex ℓ that has to be at least as much as its reported cost c_ℓ. A total budget constraint of B is assumed for payments to be made to the matched left vertices. We assume that left vertices are strategic players, which could potentially manipulate the reporting of their true cost, and hence seek a truthful algorithm, i.e., such that no incoming vertex has incentive to misreport its cost.

The considered problem is a *truthful* generalization of the online knapsack problem, where each item has a value and a weight, and the problem is to accept each item instantaneously and irrevocably, so as to maximize the sum of the values of the accepted items subject to a sum-weight constraint. If all left vertices in our problem are truthful, then keeping $p_\ell = c_\ell$ satisfies the truthfulness constraint, and the considered problem specializes to the knapsack problem with knapsack size of B. The best known algorithm for the online knapsack problem is a $10e$-competitive algorithm [1].

Copyright is held by authors. R. Vaze and M. Coupechoux are partly supported by the Indo-French CEFIPRA project "D2D for LTE-Advanced". M. Coupechoux is partly supported by the french ANR project ANR-013-INFR-004.

The considered problem is a special case of a reverse auction [2], where users submit bids for accomplishing a set of tasks and if selected, expect a payment at least as much as their reported bids. The generality is in not enforcing the one-to-one matching constraint, i.e., one user can do more than one task, or any task can be assigned to more than one user. Budget feasible mechanisms have been introduced in [3] for reverse auctions (set R is made of a single vertex). For the offline problem, with a matching constraint similar to this paper, a 3-approximate algorithm has been derived in [4] that is one-sided truthful. When the goal is to maximize the number of assigned tasks, [5] provides a 320-competitive truthful algorithm assuming the secretary-input model.

Some important applications of the considered problem are in crowdsourcing and device-to-device (D2D) communication, that is expected to be part of modern wireless communication standards. For crowdsourcing applications, a platform advertises a set of tasks it wants to accomplish and multiple users successively bid for completing those tasks and expect some payment towards that end. The job of the platform is to select the set of users that maximize its utility under a total budget constraint. If the platform is not careful in selecting its payment strategy, bidders have an incentive to misreport their costs [6, 7], and therefore we seek a truthful algorithm for matching users to tasks and decide on corresponding payments.

From a D2D communication perspective, consider a single base station (BS) and a set of users U that are connected to that BS. Let $R \subseteq U$ be the set of active users, while the remaining $L = U \backslash R$ are inactive but can potentially help users of set R to relay their communication to/from the BS, as envisaged in future networks. For any inactive user $\ell \in L$, the set of active users it can help is $R(\ell) \subseteq R$ and at any time it can only help any one user from $R(\ell)$. Since relaying requires ℓ to spend some of its resources, e.g., battery, it is natural to assume that ℓ expects some payment for its help, and submits a corresponding bid at the time of advertising its inactive state and ability to help. The job of the BS is to allocate at most one helper (matching) from set L for each user in set R, and decide the corresponding payment that is at least as much as the bid of that user in L. Clearly, there is incentive for users in set L to misreport their bids in order to extract more payments and this motivates the need to seek truthful matching algorithm.

The main contribution of this paper is a 24β-competitive randomized matching algorithm that is truthful and satisfies the payment budget constraint, where β is the ratio of the largest to the smallest utility of any edge. To keep the problem non-degenerate, similar to other prior related works on online algorithms [1, 8], we consider a secretarial input model, where the order of arrival of left vertices is uniformly random, but their utilities and bids can be arbitrary or even adversarial. Under this model, we modify the offline algorithm [4] and then use the sample and price algorithm [1] to make the algorithm online, where the novelty is in terms of defin-

ing the price for each right vertex and the corresponding payment for any left vertex that is matched to the right vertex. To contrast with the online knapsack problem, which we noted is a special case where truthfulness is guaranteed, the price of truthfulness is $24\beta/10e$ in terms of competitiveness.

2. ONLINE TRUTHFUL BUDGETED MATCHING

Let $G = (L \cup R, E)$ be a bipartite graph with left vertices L, and right vertices R, and edge set E. The utility or weight of each edge e is $u(e)$. For a set of edges E', its utility is $u(E') = \sum_{e \in E'} u(e)$. Each left vertex ℓ has an associated cost or bid c_ℓ, that does not depend on the right vertex r. We denote $c(e) = c_\ell$ as the cost/bid of edge $e = (\ell, r)$. If a left vertex ℓ is matched to a right vertex r, then a minimum payment of c_ℓ has to be made to user ℓ, with an overall budget constraint of B. Let $u_{max} = \max_e u(e)$ and $u_{min} = \min_e u(e)$. We assume that $\frac{u_{max}}{u_{min}} \leq \beta$. Moreover, we also assume the typical large market assumption [4], i.e., $\frac{u_{max}}{u^*}$ is small, where u^* is the optimal sum-utility of the matching under the budget constraint. Thus, no single user can influence the outcome significantly.

Remark 1. *As shown in [6], if bids of left vertices are used as payments, there is incentive for left vertices to misreport their bids, and consequently the mechanism is not truthful or incentive compatible. Thus, the payment strategy is non-trivial.*

In this work, we consider the online problem, where the set R of right vertices is known ahead of time, while the left vertices of set L arrive sequentially in time and reveal their edge set (and utilities) incident on the right vertices together with their bid. On each left vertex arrival, it has to be matched irrevocably to any one of the unmatched right vertex at that time, if at all. If a left vertex is matched, then the payment to be made to it is also decided at the time of its arrival that cannot be changed later.

To keep the problem non-degenerate in terms of competitive ratio, we assume that the order of arrival of left vertices is uniformly random, that is each permutation of left vertices is equally likely. As a result, the objective is to find a truthful algorithm with constant expected competitive ratio under the payment budget constraint of B. The weights (utilities), however, are allowed to be selected by an adversary.

Before considering the online scenario, we first discuss the offline (all left vertices and their edges are available non-causally) case, and define the optimal fractional solution (matching under budget constraint) to be $\mathsf{OPT}(B)$. Note that for defining OPT, we are using raw bids as payments and truthfulness is not required. We note an important property for $\mathsf{OPT}(B)$ whose proof is immediate.

Lemma 1. *For $\alpha \leq 1$, $u(\mathsf{OPT}(B)) \leq \frac{1}{\alpha} u(\mathsf{OPT}(\alpha B))$.*

For the offline scenario, we propose a THRESHOLD algorithm that is inspired (a modified version) by the UNIFORMMECHANISM algorithm [4], where the GREEDY subroutine is the usual greedy matching algorithm for a bipartite graph. We define for each edge $e = (\ell, r)$ a buck per bang $b(e) = \frac{c(e)}{u(e)}$ that represents the cost per unit utility. For any γ, let $G(\gamma)$ be the graph obtained by removing all edges $e \in E(G)$ with buck per bang $b(e) > \gamma$.

The main idea behind the THRESHOLD algorithm is to find the largest threshold γ_B, (subject to budget constraint), such that all edges whose buck per bang is more than that threshold are not considered for matching, while maintaining a (greedy) matching with large enough sum-utility.

Algorithm 1 THRESHOLD

1: **Input:** Graph G, Budget B, $m = |E(G)|$
2: **Output:** Matching M, Threshold γ_B
3: $\mathcal{A}(G) = \{\gamma : \sum_{e \in \mathsf{M}} \gamma u(e) \leq B,\ \mathsf{M} = \text{GREEDY}(G(\gamma))\}$
4: $\gamma_B = \max\{\gamma : \gamma \in \mathcal{A}(G)\}$
5: Accept all users in $\mathsf{M} = \text{GREEDY}(G(\gamma_B))$

Lemma 2. *Let M be the matching output by* THRESHOLD *algorithm with input graph G under budget constraint B. Then $u(\mathsf{M}) \geq \frac{\mathsf{OPT}(B)}{3}$.*

Proof. For convenience we suppress the dependence on B whenever its not essential to do so. Decompose the optimal fractional matching solution $\mathsf{OPT} = \{\mathsf{OPT}^+ \cup \mathsf{OPT}^-\}$, where OPT^+ contains edges of OPT that have $b(e) > \gamma_B$, and OPT^- contains edges of OPT that have $b(e) \leq \gamma_B$. Similarly, let $\mathsf{OPT}(\gamma_B)$ be the optimal fractional matching on subgraph $G(\gamma_B) \subseteq G$, where γ_B is the output threshold from the THRESHOLD algorithm with graph G. By definition of optimal matching, $u(\mathsf{OPT}^-) \leq u(\mathsf{OPT}(\gamma_B))$. Moreover, for M, the output matching from THRESHOLD algorithm with graph G, we have $u(\mathsf{M}) \geq \frac{u(\mathsf{OPT}(\gamma_B))}{2}$, since M is a greedy matching on $G(\gamma)$ (subgraph with all edges having $b(e) \leq \gamma_B$). Therefore, $u(\mathsf{M}) \geq \frac{u(\mathsf{OPT}^-)}{2}$.

All edges $e = (\ell, r) \in \mathsf{OPT}^+$, have $b(e) = \frac{c_\ell}{u(e)} > \gamma_B$. Thus, $u(\mathsf{OPT}^+) = \sum_{e=(\ell,r) \in \mathsf{OPT}^+} x(\ell) u(e) < \frac{\sum_{e=(\ell,r) \in \mathsf{OPT}^+} x(\ell) c_\ell}{\gamma_B}$, where $x(\ell)$ are fractional weights in the optimal solution. Moreover, the total budget constraint of B ($\sum_{e=(\ell,r) \in \mathsf{OPT}} x(\ell) c_\ell \leq B$) implies that $u(\mathsf{OPT}^+) < \frac{B}{\gamma_B}$. Assuming that the budget constraint is tight with the THRESHOLD algorithm ($\sum_{e \in \mathsf{M}} \gamma_B u(e) = B$), $u(\mathsf{M}) = \frac{B}{\gamma_B}$. Therefore, $u(\mathsf{OPT}^+) < u(\mathsf{M})$. Combining this with $u(\mathsf{M}) \geq \frac{u(\mathsf{OPT}^-)}{2}$, we have $u(\mathsf{M}^*) \leq 3u(\mathsf{OPT})$ as required.

If the budget constraint is not tight with the THRESHOLD algorithm, then under our assumption that the maximum utility of any edge is u_{max}, and by the definition of THRESHOLD algorithm that finds the largest feasible γ, the leftover budget $B - \sum_{e \in \mathsf{M}} \gamma_B u(e)$ is no more than $\gamma_B u_{max}$, and similar argument gives us that $u(\mathsf{OPT}) \leq (3 + o(1)) u(\mathsf{M})$. □

We next state a critical lemma for analyzing the proposed online version of THRESHOLD, ON.

Lemma 3. *Let $G = (L \cup R, E)$ and $F \subseteq G$, such that $F = (L \backslash L' \cup R, E')$, and the edge set E' is such that all edges incident on left vertices in set L' are removed simultaneously, while all edges incident on $L \backslash L'$ are retained as it is. Then*

$$u(\text{GREEDY}(G)) \geq u(\text{GREEDY}(F)).$$

Moreover

$$u(\text{GREEDY}(G(\gamma_1))) \geq u(\text{GREEDY}(G(\gamma_2)))$$

for $\gamma_1 \geq \gamma_2$, and $u(\text{GREEDY}(G(\gamma))) \geq u(\text{GREEDY}(F(\gamma)))$.

Proof. For arbitrary subgraph $F \subseteq G$, $u(\text{GREEDY}(G))$ may or may not be larger than $u(\text{GREEDY}(F))$. However, when a left vertex is removed (by deleting all edges incident to it), the proof of claim 1 follows standard procedure by showing that the weight of the edge incident on any right vertex in $\text{GREEDY}(G)$ is at least as much as in $\text{GREEDY}(F)$. Detailed proof is omitted for lack of space. For the second and third claim, note that an edge e incident on left vertex ℓ is removed in $G(\gamma)$ compared to G, if $b(e) > \gamma$ or equivalently if $u(e) < \frac{c_\ell}{\gamma}$. Recall that the cost of any edge only depends on the index of its left vertex. Hence, if edge $e = (\ell, r)$ is removed from G to obtain $G(\gamma)$, then all the edges e' incident on ℓ with utility $u(e') < u(e)$ are also removed. So essentially, edges

are removed monotonically from G to produce $G(\gamma)$. So the proofs for the second and third claim follow similarly to the first. □

The importance of Lemma 3 is in showing that THRESHOLD is solvable in polynomial time and the threshold γ_B is monotonic. We prove the two claims as follows.

Lemma 4. THRESHOLD *is solvable in polynomial time.*

Algorithm THRESHOLD involves finding a maximum in Step 4. We will show that one can use bisection to solve this maximization. We would like to note that if $u(\text{GREEDY}(G(\gamma))) \not> u(\text{GREEDY}(F(\gamma)))$, then finding this maximum is non-trivial.

Proof. From the definition of Algorithm THRESHOLD its clear that if any $\gamma \in \mathcal{A}(G)$, then $\gamma_B \geq \gamma$. Hence the key step is to show that if any $\gamma \notin \mathcal{A}(G)$, then $\gamma_B < \gamma$ which follows from the second claim of Lemma 3, that $u(\text{GREEDY}(G(\gamma_1))) \geq u(\text{GREEDY}(G(\gamma_2)))$ for $\gamma_1 \geq \gamma_2$. Therefore, if for any $\gamma \notin \mathcal{A}(G)$, then for any $\gamma' > \gamma$, $\gamma' \notin \mathcal{A}(G)$. Hence we can use bisection to find the maximum. □

The main and critical difference between the THRESHOLD and UNIFORMMECHANISM [4] algorithm is the maximization step that ensures the following monotonicity property on γ_B (Lemma 5) that allows us to make the algorithm *online* using the SAMPLEANDPRICE algorithm [8].

Lemma 5. *Let $G = (L \cup R, E)$ and $F = (L \backslash L' \cup R, E')$, where the edge set E' is such that all edges incident on left vertices in set L' are removed simultaneously, while all edges incident on $L \backslash L'$ are retained as it is. Then $\gamma_B(F) \geq \gamma_B(G)$.*

Proof. From Lemma 3,

$$u(\mathsf{M}(F(\gamma))) \leq u(\text{GREEDY}(G(\gamma))). \tag{1}$$

Let the threshold and the matching obtained by running THRESHOLD on G with budget B be $\gamma_B(G) = \gamma$, and $\mathsf{M}(G)$, respectively, where $\gamma \leq \frac{B}{u(\text{GREEDY}(G(\gamma)))}$. Now we consider $F(\gamma)$ as the input graph to the THRESHOLD with same budget constraint B. Since $\gamma \leq \frac{B}{u(\text{GREEDY}(G(\gamma)))}$, from (1), clearly, $\gamma \leq \frac{B}{u(\mathsf{M}(F(\gamma)))}$, and $\sum_{e \in \mathsf{M}(F(\gamma))} \gamma u(e) \leq B$. Therefore, $\gamma \in \mathcal{A}(F)$, which by definition of $\gamma_B(F)$ implies $\gamma_B(F) \geq \gamma$. □

We now describe our online algorithm ON, that produces the matching M_{ON} and associated payments for left vertices that are part of matching M_{ON}. The idea behind ON is as follows:

- Do not match any of the first half of left vertices (called the observation phase), and only use them to run the offline THRESHOLD algorithm and find the threshold $\gamma_{1/2}$ and the matching $\mathsf{M}_{1/2}$ with budget $B' = \frac{B}{\beta}$. Recall that $\beta > 1$, hence we are finding a conservative estimate for $\gamma_{1/2}$.
- For any right vertex $r \in \mathsf{M}_{1/2}$, set its value $v(r)$ to be the weight (utility) of the matched edge in $\mathsf{M}_{1/2}$. We are assuming that $|L_{1/2}|$ is large enough compared to the number of right vertices and all right vertices are matched by THRESHOLD in offline phase using the first half of the left vertices, i.e., $v(r) > 0,\ \forall\, r \in R$.
- In the decision phase, starting with the arrival of $\lfloor L/2 \rfloor + 1^{th}$ left vertex, delete all edges that have buck per bang $b(e)$ larger than $\gamma_{1/2}$. Among the surviving edges, match the left vertex to the right vertex with the largest weight (utility) that is higher than the value of the right vertex found from $\mathsf{M}_{1/2}$, if any.

Algorithm 2 ON Algorithm

1: **Input:** L set of left vertices/users that arrive sequentially, R set of right vertices, Budget $B' = \frac{B}{\beta}$,
2: $L_{1/2}$ = first half of left vertices L
3: Run THRESHOLD on $G_{1/2} = (L_{1/2} \cup R, E_{1/2})$ with budget B' to obtain $\gamma_{1/2} \triangleq \gamma_{B'}(G_{1/2})$ and matching $\mathsf{M}_{1/2}$
4: **for** each right vertex $r \in R$ **do**
5: Set value $v(r) := u(e)$ for $e = (\ell, r) \in \mathsf{M}_{1/2}$
6: **end for**
7: %Decision Phase
8: $\mathsf{M}_{\mathsf{ON}} = \emptyset$
9: **for** every new left vertex $\ell \in L \backslash L_{1/2}$ **do**,
10: %Pruning:
11: Delete all edges $e = (\ell, r), r \in R$ s.t. $b(e) > \gamma_{1/2}$
12: Let $e^\star = \arg\max_{e=(\ell,r), r \in R, u(e) > v(r)} u(e)$ be the largest weight (utility) edge after pruning with weight larger than the value of the corresponding right vertex.
13: Let $e^\star$ be incident on right vertex $r^\star$
14: **if** $\mathsf{M}_{\mathsf{ON}} \cup \{e^\star\}$ is a matching **then**
15: $\mathsf{M}_{\mathsf{ON}} = \mathsf{M}_{\mathsf{ON}} \cup \{e^\star\}$
16: Pay $p_\ell = \beta \gamma_{1/2} v(r^\star)$ to user ℓ
17: **else**
18: Let ℓ be permanently unmatched
19: **end if**
20: **end for**

- For each matched left vertex ℓ, pay $p = \beta \gamma_{1/2} v(r)$, where r is the right vertex to which the accepted edge from ℓ is matched.

We now compute the expected utilities of matchings $\mathsf{M}_{1/2}$ and M_{ON}, where the expectation is over the uniformly random left vertex arrival sequences.

Lemma 6. $\mathbb{E}\{u(\mathsf{M}_{1/2})\} \geq u(\mathsf{OPT}(\frac{B}{\beta}))/12$.

Proof. Let $B' = \frac{B}{\beta}$. Let $G = (L \cup R, E)$ be the full graph, while $G_{1/2} = (L_{1/2}(\sigma) \cup R, E_{1/2}(\sigma))$, be the graph consisting of only the first half of left vertices that depends on arrival sequence σ.

Since $G_{1/2} \subseteq G$, from Lemma 5, we have $\gamma_{B'}(G_{1/2}) \geq \gamma_{B'}(G)$. Thus, all the edges of $G(\gamma_{B'}(G))$ that are incident on left vertices $L_{1/2}(\sigma)$ are also present in the pruned graph $G_{1/2}(\gamma_{B'}(G_{1/2}))$. Let the greedy matching over the 'bigger' graph $G(\gamma_{B'}(G))$ be $\mathsf{M}(G)$. Let the subset of edges of $\mathsf{M}(G)$ that are incident on left vertices belonging to $L_{1/2}(\sigma)$ be $\mathsf{M}(G)_{\text{fh}}$. Let the optimal fractional matching on $G_{1/2}(\gamma_{B'}(G_{1/2}))$ be $\mathsf{OPT}(B')_{1/2}$, By definition, we have

$$u(\mathsf{OPT}(B')_{1/2}) \geq u(\mathsf{M}(G)_{\text{fh}}).$$

Since we are considering the uniformly random arrival model for left vertices, i.e., $L_{1/2}$ is obtained by sampling each left vertex of L with probability $\frac{1}{2}$, we have $\mathbb{E}\left\{u(\mathsf{M}(G)_{\text{fh}})\right\} \geq \frac{u(\mathsf{M}(G))}{2}$ and hence

$$\mathbb{E}\left\{u(\mathsf{OPT}(B')_{1/2})\right\} \geq \frac{u(\mathsf{M}(G))}{2}. \tag{2}$$

Moreover, since THRESHOLD computes a greedy matching over $G_{1/2}(\gamma_{B'}(G_{1/2}))$, we have $u(\mathsf{M}_{1/2}) \geq \frac{u(\mathsf{OPT}(B')_{1/2})}{2}$ for any realization σ. From Lemma 2, we already know that $u(\mathsf{M}(G)) \geq \frac{\mathsf{OPT}(B')}{3}$. Hence from (2), we have $\mathbb{E}\{u(\mathsf{M}_{1/2})\} \geq u(\mathsf{OPT}(\mathsf{B}'))/12$. □

Lemma 7. $\mathbb{E}\{u(\mathsf{M}_{\mathsf{ON}})\} \geq \mathbb{E}\{u(\mathsf{M}_{1/2})\}/2$.

Proof. Consider the graph $G(\gamma_{B'}(G))$, where $\gamma_{B'}$ is the output threshold of THRESHOLD algorithm with budget B' and graph $G_{1/2}$. Then the setting of value for each right vertex, and the greedy

selection of edges with weights larger than the value of right vertices in the decision phase of ON is identical to running SAMPLEANDPRICE algorithm (Algorithm 3) on graph $G(\gamma_{B'}(G))$ with $p = \frac{|L_{1/2}|}{|L|} = \frac{1}{2}$, and hence it follows from Lemma 2.5 [8] that $\mathbb{E}\{u(\mathsf{M}_{\mathsf{ON}})\} \geq \mathbb{E}\{u(\mathsf{M}_{1/2})\}/2$. □

Algorithm 3 SAMPLEADPRICE Algorithm

Input: $G = (L \cup R, E)$ and $p \in [0, 1]$
$k \leftarrow Binomial(|L|, p)$
Let L' be the first k vertices of L
$M_1 \leftarrow$ GREEDY(G'), with $G' = (L' \cup R, E')$
for each $r \in R$ **do**
    Set $v(r) = u(e)$ of the edge e incident to r in M_1
end for
$M \leftarrow \emptyset$
for each subsequent $\ell \in L \backslash L'$ **do**
    Let $e = (\ell, r)$ be the highest-weight edge with $u(e) \geq v(r)$
    if $M \cup \{e\}$ is a matching **then**
        $M = M \cup \{e\}$
    end if
end for

The following theorem is the main result of the paper.

Theorem 1. *Algorithm* ON *is* 24β*-competitive, satisfies the budget constraint and is truthful.*

Proof. The 24β-competitiveness of ON follows from combining Lemma 1,6, and 7. The budget feasibility and truthfulness are shown in Lemma 8 and 9, respectively. □

Lemma 8. *Algorithm* ON *satisfies the payment budget constraint, and payment* $p_\ell \geq c_\ell$ *for any selected left vertex* ℓ*.*

Proof. Let $\gamma_{1/2} = \gamma_{B'}(G_{1/2})$ for simplicity. For a selected left vertex ℓ, its buck per bang $\frac{c_\ell}{u(e)} \leq \gamma_{1/2}$ and $u(e) > v(r)$, where $e = (\ell, r)$ is the selected edge. From the definition of β, $\beta v(r) \geq u(e)$. Thus, $p_\ell = \gamma_{B'}\beta v(r) \geq \gamma_{1/2} u(e)$ and hence $p_\ell \geq c_\ell$.

From the definition of algorithm THRESHOLD, we know that the output threshold $\gamma_{B'} \in \mathcal{A}(G)$ for the offline phase (first half of left vertices) of ON, and hence

$$\gamma_{1/2} \sum_{e \in \mathsf{M}_{1/2}} u(e) \leq B'. \tag{3}$$

Since the value of any right vertex r is $v(r) = u(e)$ where $e = (\ell, r) \in \mathsf{M}_{1/2}$. Therefore, from (3), we have that

$$\sum_{r: e=(\ell,r) \in \mathsf{M}_{1/2}} v(r) \leq \frac{B'}{\gamma_{1/2}}. \tag{4}$$

Clearly, in the decision phase of ON, at most one left vertex is selected for each right vertex, and the payment made is $p_\ell = \beta\gamma_{B'} v(r)$ if ℓ is matched, and $p_\ell = 0$ otherwise. Thus, the total payment made

$$\sum_{\ell, e=(\ell,r) \in \mathsf{M}_{\mathsf{on}}} p_\ell \leq \sum_{r, e=(\ell,r) \in \mathsf{M}_{1/2}} \beta\gamma_{1/2} v(r).$$

Thus, from (4), we get that $\sum_{\ell, e=(\ell,r) \in \mathsf{M}_{\mathsf{on}}} p_\ell \leq B$. □

Next, we show the most important property of ON, its truthfulness. Towards that end, we will use the Myerson's Theorem [2].

Theorem 2. *[2] A reverse auction is truthful if and only if:*

- *The selection rule is monotone. If a user* ℓ *wins the auction by bidding* c_ℓ*, it would also win the auction by bidding an amount* c'_ℓ*, where* $c'_\ell < c_\ell$*.*
- *Each winner is paid a critical amount. If a winning user submits a bid greater than this critical value, it will not get selected.*

Lemma 9. ON *is a truthful online algorithm.*

Proof. As stated before, the considered problem is a special case of a reverse auction. Thus to ensure that ON is truthful, we show that both the conditions of Theorem 2 are satisfied. The monotone condition is easy to check, since in the decision phase, if any left vertex reduces its bid, then clearly its buck per bang $b(e)$ decreases, and hence it is still accepted if it was accepted before.

The second condition of payment being critical is also satisfied, shown as follows. Note that the payment made by ON to a selected left vertex ℓ is $p_\ell = \beta\gamma_{B'}(G_{1/2})v(r), e = (\ell, r) \in \mathsf{M}_{on}$, where the right vertex index r is such that utility $u(e), e = (\ell, r)$ is largest among the unmatched right vertices at the time of arrival of vertex ℓ that have an edge to left vertex ℓ, and $u(e) > v(r)$.

Recall that $\frac{u_{max}}{u_{min}} \leq \beta$, hence $\frac{u(e)}{v(r)} \leq \beta$, where $e = (\ell, r)$. Now, if suppose the bid c_ℓ of left vertex ℓ is more than $p_\ell = \beta\gamma_{B'}(G_{1/2})v(r)$, then since $u(e) \leq \beta v(r)$, buck per bang of left vertex ℓ, $c_\ell/u(e) > \gamma_{B'}(G_{1/2})$. Moreover, since $u(e) > u(e')$ for all edges e' incident on unmatched right vertices from ℓ at the arrival of left vertex ℓ, we have that $c_\ell/u(e') > \gamma_{B'}(G_{1/2})$. Thus, all edges out of left vertex ℓ incident on currently unmatched right vertices are removed in the pruning stage of the decision phase, and hence vertex ℓ cannot be selected. □

3. REFERENCES

[1] Moshe Babaioff, Nicole Immorlica, David Kempe, and Robert Kleinberg. A knapsack secretary problem with applications. In *Approximation, randomization, and combinatorial optimization. Algorithms and techniques*, pages 16–28. Springer, 2007.

[2] Roger B Myerson. Optimal auction design. *Mathematics of operations research*, 6(1):58–73, 1981.

[3] Y. Singer. Budget feasible mechanisms. In *Foundations of Computer Science (FOCS), 2010 51st Annual IEEE Symposium on*, pages 765–774, Oct 2010.

[4] Gagan Goel, Afshin Nikzad, and Adish Singla. Matching workers expertise with tasks: Incentives in heterogeneous crowdsourcing markets. In *NIPS Workshop on Crowdsourcing*, 2013.

[5] Yaron Singer and Manas Mittal. Pricing mechanisms for crowdsourcing markets. In *Proceedings of the 22Nd International Conference on World Wide Web*, WWW '13, pages 1157–1166, Republic and Canton of Geneva, Switzerland, 2013. International World Wide Web Conferences Steering Committee.

[6] Dejun Yang, Guoliang Xue, Xi Fang, and Jian Tang. Crowdsourcing to smartphones: incentive mechanism design for mobile phone sensing. In *Proceedings of the 18th annual international conference on Mobile computing and networking*, pages 173–184. ACM, 2012.

[7] Ashwin Subramanian, G Sai Kanth, Sharayu Moharir, and Rahul Vaze. Online incentive mechanism design for smartphone crowd-sourcing. In *Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks (WiOpt), 2015 13th International Symposium on*, pages 403–410. IEEE, 2015.

[8] Nitish Korula and Martin Pál. Algorithms for secretary problems on graphs and hypergraphs. In *Automata, Languages and Programming*, pages 508–520. Springer, 2009.